(19)日本国特許庁(JP)　(12)公表特許公報(A)　(11)特許出願公表番号
特表2006-528394
(P2006-528394A)
(43)公表日　平成18年12月14日(2006.12.14)

(51)Int.Cl.		F I			テーマコード（参考）
G06F 13/16	(2006.01)	G06F	13/16	510A	5B060
G06F 12/00	(2006.01)	G06F	12/00	550K	
G06F 12/06	(2006.01)	G06F	12/06	515F	

審査請求 未請求　予備審査請求 未請求　（全 23 頁）

(21)出願番号　特願2006-532883 (P2006-532883)
(86)(22)出願日　平成16年5月10日 (2004.5.10)
(85)翻訳文提出日　平成18年1月11日 (2006.1.11)
(86)国際出願番号　PCT/US2004/014441
(87)国際公開番号　WO2004/102403
(87)国際公開日　平成16年11月25日 (2004.11.25)
(31)優先権主張番号　60/470,078
(32)優先日　平成15年5月13日 (2003.5.13)
(33)優先権主張国　米国 (US)

(71)出願人　591016172
アドバンスト・マイクロ・ディバイシズ・インコーポレイテッド
ＡＤＶＡＮＣＥＤ　ＭＩＣＲＯ　ＤＥＶＩＣＥＳ　ＩＮＣＯＲＰＯＲＡＴＥＤ
アメリカ合衆国、９４０８８－３４５３　カリフォルニア州、サニィベイル、ピィ・オゥ・ボックス・３４５３、ワン・エイ・エム・ディ・プレイス、メイル・ストップ・６８（番地なし）
(74)代理人　100099324
弁理士　鈴木　正剛
(74)代理人　100111615
弁理士　佐野　良太

最終頁に続く

(54)【発明の名称】シリアルメモリインターコネクトを介して複数のメモリモジュールに接続されたホストを含むシステム

(57)【要約】
メモリモジュール（１５０Ａ～Ｂ）のシリアル接続されたチェーンに結合されたホスト（１００）を含むシステム（５０）。１つの実施形態において、メモリモジュールの各々は、メモリモジュール上の複数のメモリチップ（２６１）へのアクセスを制御するためのメモリ制御ハブ（１６０）を含む。メモリモジュールは、複数のメモリリンク（１１０）を介してホストへチェーン状にシリアル結合される。各メモリリンクは、ホストの方向へトランザクションを伝達するためのアップリンク（２１１）と、ホストで生じたトランザクションをチェーンの次のメモリモジュールに伝達するためのダウンリンク（２１２）とを含んでもよい。アップリンクおよびダウンリンクは、制御およびコンフィギュレーションパケットと、メモリアクセスパケットとを含むパケットを用いて、トランザクションを伝達してもよい。メモリ制御ハブは、トランザクションの復号とは関係なく、第２のメモリリンクの第２のダウンリンク上で、第１のメモリリンクの第１のダウンリンク上で受信したトランザクションを伝達してもよい。

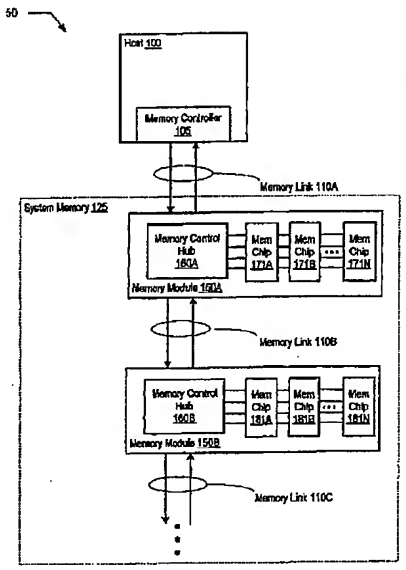

JP Publication 2006-528394

DETAILED DESCRIPTION

[Detailed Description of the Invention]
[Field of the Invention]
[0001]
This invention relates to the configuration of a memory module, and the topology of a memory subsystem in more detail about a computer system memory.
[Background of the Invention]
[0002]
The main system memory which may be constituted according to the necessity for an end user is adopted as many computer systems. In such a system, a mother board or a system board may also contain many memory expansion sockets. The one or more small circuit boards called a memory module may be inserted in a socket if needed of increasing the memory space of a computer system. Each of a memory module contains typically two or more memory devices which give the memory space of a given quantity. Generally a memory device is given using a certain type of a dynamic random access memory (DRAM:dynamic random access memory). Some DRAM type examples contain double data rate SDRAM (DDR SDRAM:double data rate SDRAM) various type with synchronous DRAM (SDRAM:synchronous DRAM).
[0003]
In the conventional computer system, a memory module is connected to a memory/DRAM controller via an address, control, and a memory bus including a data signal. The computer system which may multiplex, therefore may share the same wire set also has an address, control, and a data signal. The computer system which may use another wire also has an address, control, and a data signal. When each of an address and a control signal is inserted in the case of which, it is sent to each extended socket so that parallel connection of the memory module may be carried out to a memory/DRAM controller. If the system which may exist on the same integrated circuit (IC:integrated circuit) chip as a system processor also has a memory/DRAM controller, The system which may exist in one IC (for example, north bridge) of a chip set also has a memory/DRAM controller.
[Description of the Invention]
[Problem(s) to be Solved by the Invention]
[0004]
Although it is continuing accelerating the working speed of a computer system processor, the relative performance of the main system memory is not necessarily high at the same speed. It can become a cause that this has at least a gradual improvement of the bandwidth of the memory bus architecture selectively mentioned above.
[Means for Solving the Problem]
[0005]
Various embodiments of a system containing two or more memory modules are indicated. A host is combined with a chain with which serial connection of the memory module was carried out in one embodiment.
[0006]

In one specific embodiment, each of a memory module may also contain a memory control hub which may control access to two or more memory chips on a memory module. Serial combination of the memory module is carried out via two or more memory links at the shape of a chain at a host. Each memory link may also include uplink for transmitting a transaction to hosts, and a down-link for transmitting a transaction produced from a host to the next memory module of a chain. Each of uplink and a down-link may be a one-way link including two or more signals which may transmit a transaction using a packet containing control and a configuration packet, and a memory access packet.

[0007]

In another specific working example, a memory control hub may decode a transaction received on the 1st [of the 1st memory link] down-link. A memory control hub may transmit a transaction on the 2nd [of the 2nd memory link] down-link regardless of decoding of a transaction.

[0008]

In further another specific working example, a memory control hub may answer decoding a transaction to a memory command which has a memory address relevant to a memory control hub, and a memory address to adjust, and may also contain a DRAM controller which may access a memory chip.

[0009]

Although there is room of various corrections and another gestalt in this invention, a specific embodiment of this invention is shown in illustration in Drawings, and is indicated in detail in Description of this application. However, Drawings and a detailed description are not what was meant so that this invention might be restricted to an indicated specific gestalt, Please understand having intention of a range covering the meaning of this invention which was specified by attached Claims and all the correction things within the limits, an equivalent, and a substitute. A title is only the systematic purpose and please mind that it is not what meant using in order to restrict or interpret Description of this application or Claims. A term of "there there being also a case of - which may be - Carried out (may)", a meaning (- that is, it must carry out) of duty -- not but (must), Please care about being used through an application concerned in the sense of permission (that is, a thing which may be - done (having the potential to) and to - do is possible (being able to)). A term of "- being included (include)", and its derivative mean ["it not being limited to it, although - is included (including, but not limited to)", and]. A term of "being connected (connected)", A term of meaning "being connected directly or indirectly (directly orindirectly connected)", and "being combined (coupled)", It means "being combined directly or indirectly (directly or indirectly coupled)."

[Best Mode of Carrying Out the Invention]

[0010]

Hereafter, reference of drawing 1 shows the block diagram of one embodiment of the system containing the chain with which serial connection of the memory module was carried out. The system 50 contains the host 100 combined with the system memory 125 via the memory link 110A. The system 50 may be constituted, for example so that it may operate as some computing devices, such as a computer system or a server system. The system memory 125 contains the memory module 150A combined with the memory module 150B via the memory ink 110B. It is combined with the memory link 110C, and

the memory module 150B is shown, and as long as the memory link 110C needs to form the chain with which serial connection of the memory module combined with the host 100 was carried out, it may be combined with the further memory module (not shown). Although two memory modules are shown in the shape of a chain, please care about that it is thought that one or more memory modules may be connected in this way. It should also care about a certain thing further that the component containing the included reference number with which reference characters continue back is generally referred to only with a reference number. For example, generally, when referring to all the memory modules, the memory module 150 may be referred to.

[0011]

In the illustrated embodiment, the memory module 150A may be combined with two or more memory devices to which the memory control hubs 160A were indicated to be the memory chips 171A-171N including the memory control hub 160A, and N may be arbitrary numbers here if needed. In one embodiment, the memory control hub 160A may be combined with a memory chip via memory INTAKONEKUTO arbitrary type. For example, in one embodiment, memory INTAKONEKUTO may be a thing of a typical address, control, and the configuration of a data bus.

[0012]

Similarly the memory module 150B may be combined with two or more memory devices to which the memory control hubs 160B were indicated to be the memory chips 181A-181N including the memory control hub 160B, and N may be arbitrary numbers here if needed. In one embodiment, the memory control hub 160B may be combined with a memory chip via memory INTAKONEKUTO arbitrary type, as mentioned above. The memory chips 171A-171N, and 181A-each of 181N may be memory devices arbitrary type [, such as a memory device of a DRAM system of a memory device,], for example.

[0013]

In the illustrated embodiment, the memory links 110A-110C form memory INTAKONEKUTO. In one embodiment, each of the memory links 110A-110C forms point-to-point memory INTAKONEKUTO given as a two-set one-way line. The one-way line of one set is called a down-link, and it is constituted so that a transaction may be transmitted in the direction of a downstream from the host 100. The one-way line of another set is called uplink, and it is constituted so that a transaction may be transmitted to the host's 100 direction in an upstream direction. In one embodiment, the one-way line of each set may be given using several different signal pairs. In one embodiment, each memory link 110 is a signal pair from which each bit differs including an 18-bit down-link and 16-bit uplink. Memory INTAKONEKUTO formed by the memory link 110 may be constituted so that a packet may be transmitted, so that it may indicate still in detail below with the description of drawing 5 A - drawing 5 D.

[0014]

Speaking generally, all the transactions from the host 100, It flows into a downstream via all the memory modules 150 on a down-link, and all the response transactions flow into an upstream via each upstream memory module 150 on a rise ring from the memory module 150 which answers. As long as it says in more detail, in one embodiment, the host 100 may request that data is pulled out or stored in the system memory 125. Answering the host's 100 request, the memory controller 105 starts corresponding transactions, such as a memory read transaction or a memory write transaction, for

example. The memory controller 105 transmits a transaction to the system memory 125 via the memory link 110A. A transaction is received by the memory control hub 160A of the memory module 150A in the illustrated embodiment.
[0015]
Reception of a transaction is answered, and without decoding a transaction, the memory control hub 160A is constituted so that the transaction received to the memory module 150B via the memory link 110B may be transmitted. This is called downstream transmission of a transaction. Thus, each transaction received by the given memory control hub 160 of the given memory module 150 is transmitted to the next memory module 150 of the chain combined with the down-link, without decoding a transaction. In one embodiment, decoding of a transaction may be produced in parallel with transmission of a transaction. In other embodiments, decoding of a transaction may be produced, after a transaction is transmitted. The further detailed description about a downstream transfer function can be seen below in the description of drawing 3.
[0016]
Similarly, if the memory controller 105 starts a read request transaction, the memory module 150 which has the memory locations corresponding to the address of a request will answer with the requested data, for example. This response is on uplink of a memory module, and is transmitted to the host 100. If there is a memory module which intervenes between a transmitting memory module and the host 100, the intervening memory module will be on uplink and will transmit a response transaction to either of the next memory modules of the chain in the host 100 or an upstream direction. A response memory module may supply a response to the transaction sequence by which upstream transmission is carried out on the rise ring, if the preparation which transmits a response is completed. The further detailed description of an upstream transfer function is seen and sold to below in the description of drawing 5.
[0017]
In one embodiment, even if there is no telling which memory modules 150A and 150B it is that a specific address is related, the memory controller 105 may be constituted so that it may request to the system memory 125. For example, each of the memory module 150 may be assigned to a memory address range into a system configuration sequence. Each memory control hub 160 may also contain the logic (not shown in drawing 1) which may decode the address of a request to input. Thus, the memory control hub 160 of the given memory module 150, Decoding of the memory requests which have an address in the address range assigned to the given memory module 150 may be answered, and a memory read cycle or a memory write cycle may be started to the memory chip on the given memory module 150. In order to start a memory cycle to the memory chip connected, in one embodiment, each memory control hub 160 may also contain a DRAM controller (not shown in drawing 1), so that it may indicate still in detail below with the description of drawing 2.
[0018]
In one embodiment, the memory controller 105 may start the following memory access request, before receiving the response to a front memory access request. In such an embodiment, the memory controller 105 may process a response in a different order from what could pursue the remarkable request, therefore was transmitted.
[0019]

In another embodiment, each of the memory link 110 may form point TSUU point memory INTAKONEKUTO given as one-set bidirectional lines. Thus, the transaction of both an upstream and a downstream may flow on a bidirectional wire set. In such an embodiment, bidirectional lines may be given using several different signal pairs.
[0020]
Reference of drawing 2 shows the block diagram of one embodiment of memory modules, such as a memory module shown in drawing 1. In order to clarify and simplify, the same reference number is given to the component corresponding to what is shown in drawing 1. The memory module 150 contains the memory control hub 160 combined with the memory chips 261A-261N via the memory bus 265. The memory control hub 160 contains the control unit 240 combined with DRAM controller 250. DRAM controller 250 is combined with the memory chips 261A-261N. The control unit 240 includes the uplink control 241 and the down-link control 242. As mentioned above, the memory bus 265 may be memory INTAKONEKUTO arbitrary type. The memory control hub 160 is combined with the memory link 110A in an upstream direction, and the memory link 110B which exists in the direction of a downstream in the illustrated embodiment. The clock frequency of the memory bus 265 should mind further that it is unrelated to the clock frequency of the memory link 110.
[0021]
In the illustrated embodiment, the uplink control unit 241 may be constituted so that the packet which received from another memory module downstream may be received and transmitted. An upstream transaction sequence arises by receiving and transmitting an upstream packet. The uplink control unit 241 may be constituted so that the packet generated within the memory module 150 may be supplied to a transaction stream.
[0022]
If it may be constituted so that the packet which the down-link control unit 242 generates in a host may be received in the illustrated embodiment, and a memory module is connected to a downstream, It may be constituted so that these packets may be transmitted to a downstream memory module. The down-link control unit 242 may be constituted so that a packet may be copied and decoded. Including the address which has a packet in the address range assigned to the memory module 150 in one embodiment, if a packet is a memory access request, The down-link control unit 242 may send the command relevant to a packet to DRAM controller 250. In one embodiment, DRAM controller 250 answers the memory command from the memory control hub 160, and it is constituted so that a memory cycle may be started to the memory chips 261A-261N. However, as long as a packet is not memory requests but a configuration packet, the down-link control unit 242 may send the configuration command relevant to a packet to the core logic of the control unit 240 (not shown) for processing. If a packet does not include the address in the address range assigned to the memory module 150 in one embodiment, the memory control hub 160, If the memory module 150 is a memory module of the last of a chain, please care about that a packet may be canceled or canceled.
[0023]
In one embodiment, the memory control hub 160, When it is constituted so that a module existence signal (not shown) may be received, and this signal is activated with a downstream memory module, it is shown in an upstream memory module that a downstream memory module exists. In such an embodiment, as long as it is determined

that there is no downstream memory module in which the memory control hub 160 receives a transaction, and exists, the memory control hub 160 may cancel a transaction. In one specific working example, if the given memory module 150 is inserted in a socket, the inserted memory module may apply the signal ground to a module existence signal. Thus, an active module existence signal is an active low signal.

[0024]

When the signal ground is applied in the embodiment mentioned above, a module existence signal is active, but in other embodiments, please care about that it is possible that other voltage levels may be applied to a module existence signal so that it is shown that a memory module exists.

[0025]

Reference of drawing 3 shows the block diagram of one embodiment of a down-link control unit. In the embodiment shown in drawing 3, the down-link control unit 342 may express the down-link control unit 242 shown in drawing 2. The down-link control unit 342 receives a downstream transaction on the down-link 312A, and it is combined so that these downstream transactions may be transmitted on the down-link 312B. In one embodiment, the down-links 312A and 312B may express each of the down-links 212A and 212B of drawing 2. As mentioned above, the down-links 312A and 312B should care about that two or more bits are included. For example, in one working example, each of the down-links 312A and 312B may be an 18-bit down-link. The down-link control unit 342 receives the busy signal 371 from an upstream memory module or the host 100, and it is constituted so that the busy signal 372 may be transmitted to a downstream memory module.

[0026]

In the illustrated embodiment, the down-link control unit 342 contains the phase alignment unit 310 which receives an input transaction. The phase alignment unit 310 is combined with the transmission unit 315 and the data restoration unit 320. The data restoration unit 320 is combined with the synchronous first in first out buffer (FIFO:firstin first out) 325 combined with the address decoding logic 330. Synchronous FIFO325 is combined with DRAM controllers, such as DRAM controller 250, and the core logic in the control unit 240 of drawing 2 in one embodiment. In the embodiment shown in drawing 3, the phase alignment unit 310, the transmission unit 315, and the data restoration unit 320 should care about that it may operate in each bit of a down-link independently by bitwise. In other embodiments, the phase alignment unit 310, the transmission unit 315, and the data restoration unit 320 can consider that it may operate in all the bits of a down-link simultaneously substantially.

[0027]

In order to give required memory band width, the memory control hub 160 can become important [transmitting the received transaction to a downstream memory module efficiently]. Therefore, the transaction which the down-link control unit 342 received is thoughtlessly transmitted to a downstream. In one embodiment, phase alignment is carried out by the phase alignment unit 310, and each receiving bit is transmitted to the transmission unit 315, in order to transmit by the down-link 312B, without decoding. Each receiving bit is transmitted to the data restoration unit 320, in order to process.

[0028]

In one embodiment, in order that the down-link control unit 342 may enable it to sample

in the center of a received-data eye, the phase alignment unit 310 is constituted to each bit so that the phase of local sampling clocks may be adjusted dynamically. For example, in one embodiment, a reference clock is given from a host to the phase locked loop (PLL:phased locked loop) (not shown) in the down-link control unit 342. PLL generates a transmission clock with local sampling clocks in the down-link control unit 342. The phase alignment unit 310 is constituted so that tracking of the average phase of an input data stream may be carried out. The phase alignment unit 310 compensates again the static skew produced by a process variation, Tracking of the low frequency wave change of the data phase produced by change of voltage and temperature may be carried out, and it may exist in a reference clock, and it may be constituted so that the arbitrary low frequency wave phase jitters in which tracking is not carried out by PLL of a hub may be compensated. The phase alignment unit 310 filters the high frequency jitter produced by intersymbol interference, a cross talk or the reflection noise, and the high frequency phase noise that may be generated with a transmitter.

[0029]

In order to perform phase alignment of local sampling clocks, it may be required to give a sufficient number for each bit of the down-link 312A of transition (or penetration density). In one embodiment, transition density will be given by descrambling received data, if it receives after carrying out the scramble of the data transmitted to each bit. In one embodiment, the scramble of the data is carried out by carrying out EXCLUSIVE OR operation (XOR:Exclusive-OR) of the data by a pseudo-random binary sequence (PRBS:pseudo random binary sequence). A linear feedback shift register (LFSR:linear feedback shift register) may generate PRBS using a given seed polynomial. In another embodiment, transition density may be given by transmitting the synchronous packet which has predetermined data payloads. In such an embodiment, predetermined data payloads may also contain random or the pseudo-random pattern which has a sufficient number of transition.

[0030]

It is not concerned with the acquiring method of transition density, but the data restoration unit 320 is constituted so that a data bit may be recovered. In the embodiment using scramble, using the same PRBS as what is used in order to carry out the scramble of the data, the data restoration unit 320 may be constituted so that received data may be descrambled. On the contrary, in order to gain transition density, in the embodiment to be used, a synchronous packet the data restoration unit 320, After the phase alignment unit 310 ends alignment of local sampling clocks, only canceling synchronous packet data or canceling them also has it. [good]

[0031]

In the illustrated embodiment, synchronous FIFO325 is constituted so that the input data bit for using it by the core logic of the control unit 240 may be stored. for example, a transaction -- a packet -- since it is-izing and transmitted, a receiving bit is stored in synchronous FIFO325, and it may be reformatted in order to decode. Subsequently, a receive packet is decoded by the address decoding logic 330. As mentioned above, a memory address range may be assigned to each memory module 150. The hub address for using it for each memory control hub 160 in a configuration transaction may be assigned. Command information may be extracted and processed, as long as a transaction address is decoded and consistency is shown. In one embodiment, as long as a transaction

is a memory access request which has an address which adjusts the memory address relevant to the memory control hub 160, the decoded command may be transmitted to DRAM controller 250. As long as a transaction is a configuration transaction, the decoded command may be transmitted to the core logic of the control unit 240. A packet may be canceled or canceled if an address decoding logic section or the address decoding logic 330 does not detect consistency.

[0032]

In the illustrated embodiment, the phase alignment unit 335 is combined so that the input busy signal 371 may be received. The phase alignment unit 335 is combined with the data restoration unit 350 again, and this unit 350 is combined with the busy pulse injection unit 355. The busy pulse injection unit 355 is combined with the transmission unit 360.

[0033]

In the illustrated embodiment, the phase alignment unit 335 and the data restoration unit 350 are similar to each of the phase alignment unit 310 and the data restoration unit 320, and operate. However, in one embodiment, the busy signal 371 is a single bit difference part signal, and is not a multibit signal. The busy signal 371 and the busy signal 372 are used in order to leave sufficient idle time for the hub which has a hub near the bottom product of the serial chain of a memory module near the top of a serial chain to supply a packet. In one embodiment, the busy signal 371 includes a number of a packet of directions supplied to an upstream transaction sequence with all the memory modules in an upstream. Therefore, the busy signal 372 includes the total of the packet supplied to an upstream transaction sequence with all the memory modules in an upstream, and a number of a packet of directions which are locally supplied by the memory control hub 160.

[0034]

In one embodiment, the busy pulse injection unit 355 is constituted so that the data corresponding to the packet number which will be supplied to an upstream transaction sequence with all the memory modules in an upstream may be received. The busy pulse injection unit 355 is constituted so that directions may be received from the core logic of the control unit 240 corresponding to the number of the packets which will be locally supplied to an upstream transaction sequence by the memory control hub 160. Therefore, the busy pulse injection unit 355 is constituted so that the busy signal 372 may be made. In one embodiment, the busy signals 371 and 372 may also include a pulse to each injection packet. Thus, signals, such as the busy signal 372 showing five packets supplied, will include five pulses, for example. However, please take into consideration that arbitrary numbers of pulses may be included to each injection packet in other embodiments. A packet may be transmitted as a short packet or a long packet so that it may be indicated further below. Therefore, in one embodiment, in the case of the supplied long packet, as for the busy pulse injection unit 355, each merit may be, and it may also include two pulses to a packet.

[0035]

The busy pulse injection unit 355 is constituted so that the directions corresponding to the number of the busy pulses received from the upstream to the uplink control unit 241 of drawing 2 may be given. If a memory module is a module of the last of the chain by which serial connection was carried out in one embodiment, the uplink control unit 241, According to the number of the busy pulses received on the busy signal 371, one or more

NOP packets may be made to be supplied to an upstream transaction sequence. For example, in one working example, the uplink control unit 241 may be made to throw one NOP packet into an upstream transaction sequence every two busy pulses which received. However, other embodiments for which other ratios of an NOP packet and a receiving busy pulse may be used can be considered.

[0036]

In the illustrated embodiment, the transmission unit 360 is constituted so that a busy signal may be received from the busy pulse injection unit 355. In one embodiment, as mentioned above, LFSR is used for the transmission unit 360 before transmission of the busy signal 372, and it is further constituted so that the scramble of the data may be carried out. As other gestalten, the transmission unit 360 may supply a synchronous packet, as mentioned above.

[0037]

Reference of drawing 4 shows the block diagram of one embodiment of an uplink control unit. In the embodiment shown in drawing 4, the uplink control unit 441 may express the uplink control unit 241 shown in drawing 2. The uplink control unit 441 receives an upstream transaction on the uplink 411B, and it is combined so that an upstream transaction may be transmitted on the uplink 411A. The uplink control unit 241 is constituted so that a packet may be supplied to the upstream transaction sequence received on the uplink 411B. In one embodiment, the uplinks 411A and 411B may express each of the uplinks 211A and 211B of drawing 2. As mentioned above, the uplinks 411A and 411B should care about that two or more bits are included. For example, in one working example, each of the uplinks 411A and 411B may be 16-bit uplink.

[0038]

In the illustrated embodiment, the uplink control unit 441 contains the phase alignment unit 410 which receives an input upstream transaction. The phase alignment unit 410 is combined with the data restoration unit 415.

[0039]

In order to give required memory band width, the memory control hub 160 has that it is also important to transmit a receiving upstream transaction to an upstream memory module or a host efficiently. The transaction which the uplink control unit 441 received is transmitted to an upstream like the downstream control unit 342. However, the uplink control unit 441 may supply the packet generated locally to an upstream transaction sequence by contrast [the downstream control unit 342].

[0040]

In the illustrated embodiment, the phase alignment unit 410 operates like the phase alignment unit 310. Therefore, to each bit, in order that the uplink control unit 441 may enable it to sample in the center of a received-data eye as mentioned above, the phase alignment unit 410 is constituted so that the phase of local sampling clocks may be adjusted dynamically. Similarly, as mentioned above with the description of the data restoration unit 320, the data restoration unit 415 is constituted so that a data bit may be recovered.

[0041]

The data restoration unit 415 is combined with the multiplexer 430 combined with the transmission unit 435 in the illustrated embodiment. If there is no packet supplied to an

uplink transaction sequence in the uplink control unit 441, the uplink control unit 441 is constituted so that the transaction received on the uplink 411A may be transmitted. The packet boundary track unit 450 may give a control signal to the multiplexer 430, and carries out the scramble of the received transaction by the multiplexer 430, and it may enable it to transmit with the transmission unit 435 in one embodiment. In one embodiment, the transmission unit 435 may carry out the scramble of the data transmitted using LFSR, as mentioned above, but in other embodiments, the transmission unit 435 may supply a synchronous packet, as mentioned above.

[0042]

However, in one embodiment, when the packet supplied is received and stored in injection FIFO425 from hub core logic, the received data from the data restoration unit 415 are stored in maintenance FIFO420, in order to transmit later. Completion of transmission of the packet transmitted now will give the data stored in injection FIFO425 to the input of the multiplexer 430. The packet boundary track unit 450 carries out tracking of the packet boundary of the packet which received from the uplink 411B. The supplied packet with the packet boundary track unit 450, By giving a control signal to the multiplexer 430 at exact time, by not being in agreement with the transmitted packet, the scramble of the waiting packet poured in can be carried out, and it can transmit now with the transmission unit 435. If it assumes that there are a packet which is the injection schedule stored in injection FIFO425, and a packet of the transfer schedule stored in maintenance FIFO420, The packet boundary track unit 450 may control the multiplexer 430 selectively according to a fairness algorithm (fairness algorithm) to send one type of the packets to the transmission unit 435. Fairness algorithms may be arbitrary algorithms constituted so that either the injection packet to arbitrary given memory modules or forward packets might not run short. Please care about that it may be canceled instead of the NOP packet which received on the uplink 411B being stored in the maintenance FIF420, or being transmitted.

[0043]

In another embodiment, in order to improve the waiting time of the transmitted transaction traffic, the phase alignment unit 410 should care about that it may be combined with maintenance FIFO420 and the multiplexer 430 so that a dashed line shows. According to such an embodiment, waiting time may be improved when injection traffic does not exist.

[0044]

In one embodiment, before communication between the host 100 and each memory control hub 160 of each memory module 150, A start sequence may be performed in order to synchronize arbitrary scramble / descrambling logic which may exist in the host 100 and each memory control hub 160, synchronous FIFO, and phase alignment logic. For example, in one embodiment, a start sequence may also contain during reset the host 100 who transmits one or more alignment patterns to a downstream. By each memory control hub 160, an alignment pattern may be transmitted to an upstream and may be returned to the host 100. An alignment pattern may be used so that a receiver including scramble / descrambling logic can lock a synchronization.

[0045]

(Memory INTAKONEKUTO)

When drawing 1 and drawing 2 are referred to collectively, memory INTAKONEKUTO,

Including one or more high-speed point TSUU point memory links, such as the memory links 110A-110C, each of the memory links 110A-110C includes down-links, such as uplink of the uplink 211A etc., and the down-link 212A, for example. As mentioned above, in one embodiment, a down-link may be an 18-bit link and uplink may be a 16-bit link. Thus, an 18-bit down-link may also include 16 control address information (CAD:control, address and data) signal, a busy signal, and a control (CTL:Control) signal. Given uplink may also include 16 control address information (CAD) signal. However, in another embodiment, it is also possible that uplink of the uplink 211A etc. may include a CTL signal.

[0046]

Other signals may be given to each memory module 150 other than a high-speed link. For example, in one embodiment, a reset signal, a power OK signal, and a reference clock may be given to each memory module 150 from the host 100. Other signals may be given between each memory module. For example, as mentioned above, the following memory module existence signal may be given between memory modules.

[0047]

Speaking generally, classifying into a configuration and a control transaction, and a memory transaction the type of the transaction transmitted on the memory link 110. In one embodiment, since the memory control hub 160 is constituted, a configuration and a control transaction may be used. For example, a configuration and a control transaction may be used so that configuration registers may be accessed, and a memory address range may be assigned to a memory module or a hub address may be assigned to a memory module hub. A memory transaction may be used in order to access the memory locations in the memory chip 261A - 261N.

[0048]

Therefore, in one embodiment, there are two types of addressing currently supported, hub addressing and memory addressing. If hub addressing is used, 8 hub bit will identify the specific memory control hub accessed. In one embodiment, the hub address of FFh may express the simultaneous transmissive communication to all the memory control hubs. Using memory addressing, each hub decodes a lower part, in order to decode the upper part of an address bit in order to determine which hub should receive a request, and to opt for the memory locations accessed. In one embodiment, although there are 40 address bits, it is possible that the address bit of other numbers may be used if needed.

[0049]

In one embodiment, each of a memory link is constituted so that a transaction may be transmitted using one or more packets. A packet may also contain data payloads according to the type of a command with which a packet conveys those each including control and a configuration packet, and a memory access packet. Thus, the wire set which constitutes the memory link 110 may be used in order to transmit control, an address, and data.

[0050]

A packet may generally be characterized by each packet for the bit portion of a large number which transmit the information on a single bit to be included. Each packet is divided into several bit time, and all the bit portions of a packet are sampled between given bit times. Thus, control information and data share the same wire (for example, CAD wire) of a given link. In one embodiment, a packet is a multiple of a bit pair and the

bit time of the beginning of all the packets is sampled in an equal bit time so that it may indicate still in detail below. A packet starts with the control header whose length may be either a 1-bit pair or a 2-bit pair. In one embodiment, the first 5 bits of a control header are command code. Following Table 1 indicates the command code relevant to them to be a packet various type. However, the actual code shown in the 1st row is a thing for the purpose of illustration, and please care about that other codes may be used for a command given [each].

The Table 1 packet type and command code

| Direction [of code | header | command | description |] | normal response | Adlai |
| | length (BI | | | | | SUTAI |)
| | TTOTA | | | | | PU |
| | Yim | | | | | |
| It is |---|-NOP-|-operation-less/| bidirectional |- |- | 00h.
			idol-like			
			voice			
04h	2	AddrSet	ADORESUSE	down	Addr	hub
			TTO		Ack	
05h	2	AddrAck	address private seal	rise	-	-
			**			
06h	2	Ack	recognition	rise	-	-
Don't do 07h	2	Nak	recognition of, but it is /	rise	-	-	.
			error			
08h	2	SRdResp	-- short lead	rise	-	-
			response			
09h	2	LRdResp	-- long lead	rise	-	-
			response			
0Ah	2	ConfigRd	KONFIGI	down	RdResp	hub
			YURESHO			
			NRIDO			
0Ch	2	ConfigWr	KONFIGI	down	Ack	hub
			YURESHO			
			NRAITO			
0Eh	2	DIMMCtl	DIMM system	down	Ack	hub
			--			
10h	4	SMemRd	-- short memory	down	RdResp/Ack	memory
			lead			
11h	4	LMemRd	-- long memory	down	RdResp	memory
			lead			
12h	4	BlkMemWr	BUROKKUME	down	Ack	memory
			MORIRAITO			
13h	4	SbytMemWr	-- short byte	down	Ack	memory
			memory rye			
			TO			
14h	4	LbytMemWr	-- long byte	down	Ack	memory
			memory rye			
			TO			

[0051]
In one embodiment, a packet (except for an NOP packet) is transmitted with an error detection code (EDC:error detecting code). In one embodiment, although EDC is a 32-bit cyclic redundancy code (CRC:cyclic redundancy code), please care about that the thing of other EDC may be adopted in other embodiments if needed. Data is first transmitted by a least significant byte to an address being first transmitted to high-speed decoding in the memory control hub 160 in a most significant bit time. However, please care about that an address may be first transmitted in a least significant bit time, and data may be first transmitted by the most significant byte in other embodiments. A packet may convey the pay load of byte enabling and/or data. A packet without a pay load is called a header-only packet. In one embodiment, the size of the short data of a lead may be a half of the cash linesize programmed at the maximum. The size of the long data of a lead and block write may be the cash linesize programmed at the maximum. The data size of a byte light may be the greatest 64 bytes irrespective of setting out of cash linesize.
[0052]
The CTL signal other than a control header and command code information included in a packet may be used in order to transmit the information about each packet. As shown in following Table 2, some illustration CTL coding is shown.
Table 2 CTL coding of downstream use
Contents [of ** even number ** odd number ** CAD] **
** 0 ** 0 ** data or by TOINEBU**
** ** ** RUPEI load **
** 1 ** 1 ** control header **
** of a packet with a ** 0 ** 1 ** pay load
** ** **CRC **
C** of a ** 1 ** 0 ** head-only packet
** ** **RC **
[0053]
Sufficient information which enables it to insert a header-only packet into the pay load of another packet because the header of a packet differs from the value of CTL of a payload part may be given. While the light packet is transmitted on the link, this is enabling it to emit a read command, and in order to reduce the waiting time of a read command, it has a useful thing. Table 3 shows the illustration packet which contains a pay load by a tabular format. The packet of Table 3 shows that the header-only packet is inserted in a pay load between the bit times 4-7.
Table 3 Packet with pay load, and header-only packet inserted into pay load
** bit time **CTL **CAD **
1 bit of ** 0 ** 1 ** headers [15:0] **
1 bit of ** 1 ** 1 ** headers [31:16] **
** 2 ** 0 ** data bit [15:0] **
** 3 ** 0 ** data bit [31:16] **
2 bits of ** 4 ** 1 ** headers [15:0] **
2 bits of ** 5 ** 1 ** headers [31:16] **
2 bits of ** 6 ** 1 ** CRC [15:0] **
2 bits of ** 7 ** 0 ** CRC [31:16] **
** 8 ** 0 ** data bit [47:32] **

** 9 ** 0 ** data bit [64:48] **
1 bit of ** 10 ** 0 ** CRC [15:0] **
1 bit of ** 11 ** 1 ** CRC [31:16] **
[0054]
Drawing 5 A - drawing 5 D show the illustration packet which may be transmitted on the memory link 110A of drawing 1 - 110C. Hereafter, reference of drawing 5 A shows the figure of one embodiment of a configuration lead packet. In the illustrated embodiment, the configuration lead packets 510 are 16 bit width, and contain a 4-bit time or a 2-bit pair.
[0055]
5-bit command code (for example, 0Ah) is transmitted to the bit positions 0-4 between the bit times 0. Let the bit positions 5-7 be reserves. An 8-bit tag is transmitted to the bit positions 8-15. In one embodiment, since a remarkable downstream request and an upstream response packet are adjusted for example, tag values may be used by the host 100. If not used, a tag field may be set as the value of 00 h. An 8-bit hub address is transmitted to the bit positions 0-7 between the bit times 1. The number of 8-bit configuration registers is transmitted to the bit positions 8-15. The bits 0-15 of CRC, and 16-31 are transmitted to the bit positions 0-15 between the bit times 2 and 3, respectively.
[0056]
Reference of drawing 5 B shows the figure of one embodiment of a configuration light packet. In the illustrated embodiment, the configuration light packets 515 are 16 bit width, and contain a 6-bit time or a triplet pair.
[0057]
5-bit command code (for example, 0Ch) is transmitted to the bit positions 0-4 between the bit times 0. Let the bit positions 5-7 be reserves. An 8-bit tag is transmitted to the bit positions 8-15. An 8-bit hub address is transmitted to the bit positions 0-7 between the bit times 1. The number of 8-bit configuration registers is transmitted to the bit positions 8-15. The bits 0-15 of data payloads, and 16-31 are transmitted to the bit positions 0-15 between the bit times 2 and 3. The bits 0-15 of CRC, and 16-31 are transmitted to the bit positions 0-15 between the bit times 4 and 5, respectively.
[0058]
Reference of drawing 5 C shows the figure of one embodiment of a memory read packet. In the illustrated embodiment, the memory read packets 520 are 16 bit width, and contain a 6-bit time or a triplet pair. 5-bit command code (for example, 10h or 11h) is transmitted to the bit positions 0-4 between the bit times 0. Let the bit positions 5-7 be reserves. An 8-bit tag is transmitted to the bit positions 8-15.
[0059]
The length of the data which should return to the bit positions 0-5 is transmitted between the bit times 1. In one embodiment, it is shown that the value of 00 h does not have data, the value of 01 h shows the data of a 2-bit pair, and the value of 02 h shows the data of a 4-bit pair, and is the same as that of the following. A recognition packet (Ack:acknowledge packet) is returned to a requester as the length of a lead is zero. In one embodiment, since short RdResp arises with the lead of the cache line below half, either single long RdResp or two short RdResp(s) may arise with the lead of a bigger cache line than a half. The size of a cache line may be programmed by the configuration registers of the host 100 and each memory control hub 160 with software. Let the bits 6-7 be reserves.

The address bits 39-32 of the place where DRAM was requested to the bit positions 8-15 are transmitted.

[0060]

Between the bit times 2, the address bits 31-16 of the place where DRAM was requested to the bit positions 0-15 are transmitted, and the address bits 3-15 of the place where DRAM was requested in the bit positions 3-15 are transmitted between the bit times 3. A packet priority is transmitted to the bit positions 0-1 between the bit times 3. In one embodiment, a priority may express the priority of the packet to other requests. For example, even if one priority is already in progress, it may be for delaying all the requests with a lower priority, and performing this request before them. Let the bit position 2 be a reserve. The bits 0-15 of CRC, and 16-31 are transmitted to the bit positions 0-15 between the bit times 4 and 5, respectively.

[0061]

Reference of drawing 5 D shows the figure of one embodiment of a block memory write packet. In the illustrated embodiment, the block memory write circuits 525 are 16 bit width, and contain an 8-bit time or a 4-bit pair. 5-bit command code (for example, 12h) is transmitted to the bit positions 0-4 between the bit times 0. Let the bit positions 5-7 be reserves. An 8-bit tag is transmitted to the bit positions 8-15.

[0062]

The length of the data currently transmitted to data payloads is transmitted to the bit positions 0-5 between the bit times 1. In one embodiment, it is shown that the value of 00 h does not have data, the value of 01 h shows the data of a 2-bit pair, and the value of 02 h shows the data of a 4-bit pair, and is the same as that of the following. Let the bits 6-7 be reserves. The address bits 39-32 of the place of DRAM currently written in the bit positions 8-15 are transmitted.

[0063]

The address bits 31-16 of the place of DRAM currently written in the bit positions 0-15 are transmitted between the bit times 2, and the address bits 3-15 of the place of DRAM currently written in the bit positions 3-15 are transmitted between the bit times 3. A packet priority is transmitted to the bit positions 0-1 between the bit times 3. Let the bit position 2 be a reserve.

[0064]

The bits 0-15 of the 1st bit pair of data payloads, and 16-31 are transmitted to the bit positions 0-15 between the bit times 4 and 5. As long as more data is written in, the bit pair which follows may transmit the bits 0-15 of the data payloads which follow, and 16-31. The bits 0-15 of CRC, and 16-31 are transmitted to the bit positions 0-15 between bit time 4+2N and 5+2N, respectively.

[0065]

Although only the packet of four types is shown, please care about that the packet of other types which may be the things corresponding to the command code mentioned to Table 3 is also considered. Although various fields of the illustration packet which has a specific number of bits are shown, please care about further that it is also possible that various fields may contain the bit of other numbers if needed in each packet in other embodiments.

[0066]

Drawing 6 is a block diagram of one embodiment of a computer system. The computer

system 600 contains the process nodes 612A-612D to which interconnection of each was carried out by coherent packet interface link 615 A-D. Each link of the coherent packet interface 615 may form a high-speed point TSUU point link. Each of process node 612 A-D may also include one or more processes. The computer system 600 contains again the I/O node 620 combined with the process node 612A via the un-coherent packet interface 650A. The I/O node 620 may be connected to another I/O node (not shown) in chain topology by the un-coherent packet interface 650B, for example. The process node 612A is shown as a host node, and may also include the host bridge for communicating with the I/O node 620 via the NC packet interface 650A. Process node 612 B-D may also include the host bridge for communicating with other I/O nodes (not shown) again. The un-coherent packet interface link formed of NC packet interface 650 A-B may be called a point TSUU point link. The I/O node 620 is connected to the pair of peripheral-bus 625 A-B.

[0067]

Drawing 6 shows the corresponding system memory (for example, 617A and 617B) combined with the process nodes 612A and 612B. In the illustrated embodiment, each of the process nodes 612A and 612B expresses a host as shows drawing 1, and each system memory 617 may be given to the configuration described with the description of drawing 1 and drawing 2 which were mentioned above. INTAKONEKUTO between each of the process nodes 612A and 612B and the system memory 617 corresponding to them may express memory INTAKONEKUTO including the memory link 110C mentioned above in drawing 1 and drawing 2. Please care about that the process node of other numbers may be used in other embodiments. It is possible that each of the process nodes 612C and 612D may be connected like corresponding system memories, such as the system memory 617, for example.

[0068]

In the illustrated embodiment, each link of the coherent packet interface 615 is given as a set of a one-way line (for example, the line 615A). It is used in order to transmit a packet to the process node 612B from the process node 612A, and the line 615B is used in order to transmit a packet to the process node 612C from the process node 612B. Other line set 615 C-D is used in order to transmit a packet among other process nodes as shown in drawing 1. In order that the coherent packet interface 615 may communicate between process nodes, it may operate at a cash coherent ceremony ("coherent link"). In order that the un-coherent packet interface 650 may communicate between I/O nodes and among host bridges, such as a host bridge of an I/O node and the process node 612A, it may operate at an un-coherent ceremony ("un-coherent link"). The interconnection of two or more nodes through a coherent link may be called a "coherent fabric." Similarly, the interconnection of two or more nodes through an un-coherent link may be called an "un-coherent fabric." Please care about that the packet transmitted to another thing from one process node may pass one or more intermediate nodes. For example, the packet transmitted to the process node 612C from the process node 612A may pass either the process node 612B or the process node 612D, as shown in drawing 6. Arbitrary suitable routing algorithms may be used. Other embodiments of the computer system 600 may also contain the thing which made the number of process nodes fluctuate from the embodiment shown in drawing 6.

[0069]

One example of the packet interface of the un-coherent packet interface 650 etc. may suit with HyperTransport™ art. The peripheral buses 625A and 625B illustrate general peripheral buses, such as a circumference component INTAKONEKUTO (PCI:peripheral component interconnect) bus. However, please understand that the bus of other types may be used.

[0070]

Please care about further other computer system configurations being possible and thinking. For example, it is possible that the system memory configuration mentioned above in drawing 1 - drawing 5 may be used with the computer system which adopted the processor chip set containing a north bridge. In such an embodiment, the memory controller in a north bridge may function as a host.

[0071]

Although some embodiments have been mentioned above in detail, if it is a person skilled in the art, many modifications and examples of correction will become clear by recognizing the above-mentioned indication thoroughly. It has intention of being interpreted as Claims including such all the modifications and examples of correction.

[Industrial applicability]

[0072]

This invention can be applied to a computer system memory, and is generally obtained.

[Brief Description of the Drawings]

[0073]

[Drawing 1]It is a block diagram of one embodiment of the system containing the chain with which serial connection of the memory module was carried out.

[Drawing 2]It is a block diagram of one embodiment of memory modules, such as a memory module shown in drawing 1.

[Drawing 3]It is a block diagram of one embodiment of down-link control units, such as a down-link control unit shown in drawing 2.

[Drawing 4]It is a block diagram of one embodiment of uplink control units, such as an uplink control unit shown in drawing 2.

[Drawing 5 A] It is a figure of one embodiment of a configuration lead packet.

[Drawing 5 B] It is a figure of one embodiment of a configuration light packet.

[Drawing 5 C] It is a figure of one embodiment of a memory read packet.

[Drawing 5 D] It is a figure of one embodiment of a memory write packet.

[Drawing 6]It is a block diagram of one embodiment of a computer system.

【００７２】
本発明は、一般的に、コンピュータシステムメモリに応用可能でありうる。
【図面の簡単な説明】
【００７３】
【図１】メモリモジュールのシリアル接続されたチェーンを含むシステムの１つの実施形態のブロック図である。
【図２】図１に示すメモリモジュールなどのメモリモジュールの１つの実施形態のブロック図である。
【図３】図２に示すダウンリンク制御ユニットなどのダウンリンク制御ユニットの１つの実施形態のブロック図である。
【図４】図２に示すアップリンク制御ユニットなどのアップリンク制御ユニットの１つの実施形態のブロック図である。
【図５Ａ】コンフィギュレーションリードパケットの１つの実施形態の図である。
【図５Ｂ】コンフィギュレーションライトパケットの１つの実施形態の図である。
【図５Ｃ】メモリリードパケットの１つの実施形態の図である。
【図５Ｄ】メモリライトパケットの１つの実施形態の図である。
【図６】コンピュータシステムの１つの実施形態のブロック図である。

【図１】

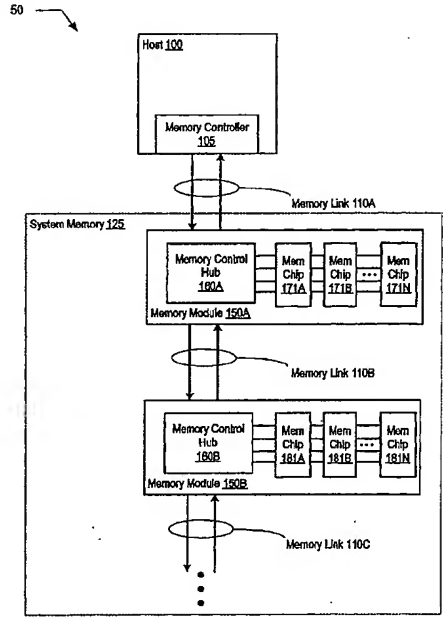

FIG. 1

【図２】

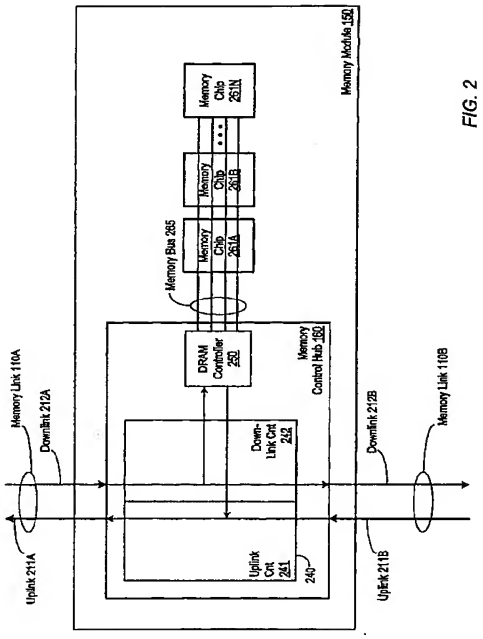

FIG. 2

【図３】

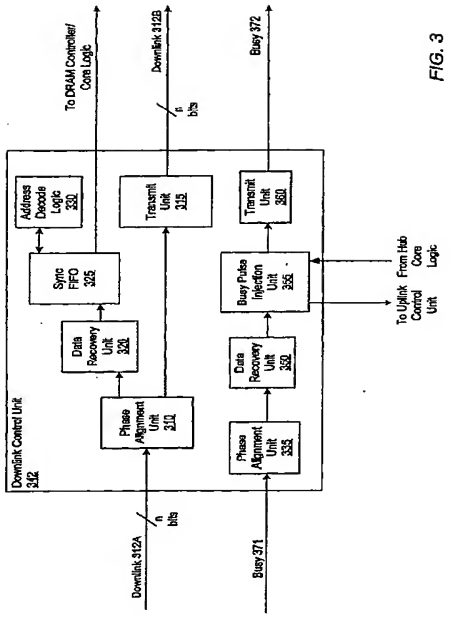

FIG. 3

【図４】

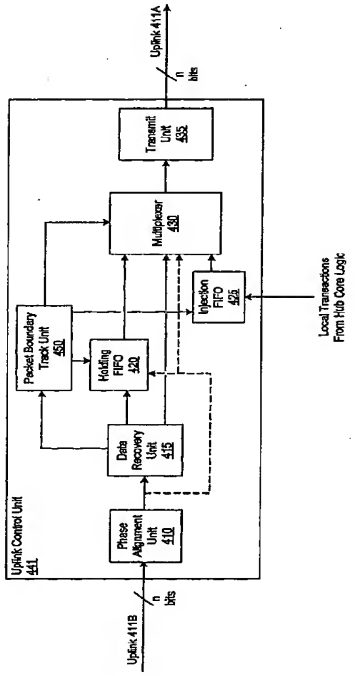

FIG. 4

【図５Ａ】

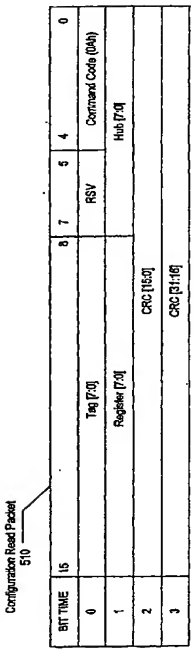

FIG. 5A

【図５Ｂ】

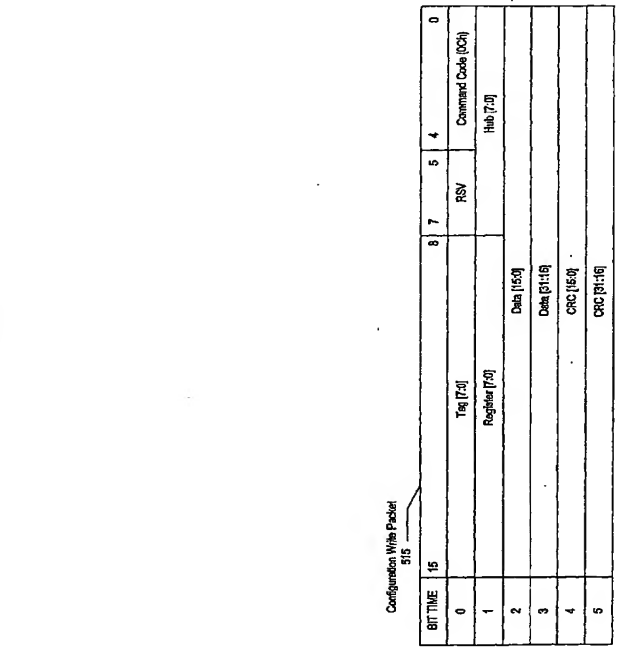

FIG. 5B

【図５Ｃ】

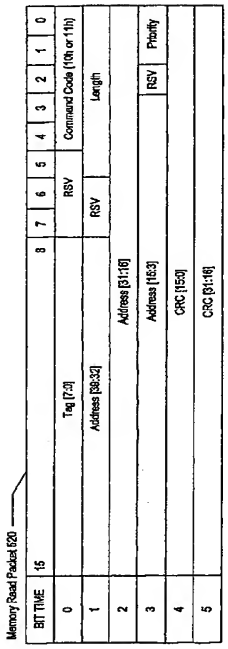

FIG. 5C

【図５Ｄ】

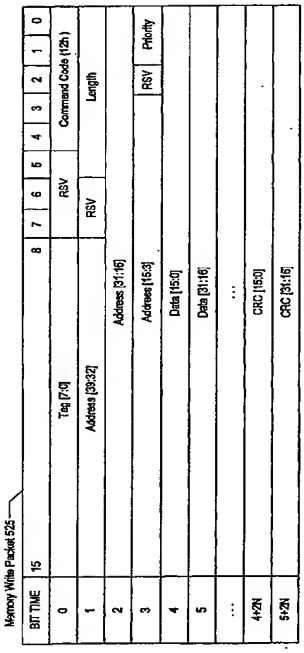

FIG. 5D

【図６】

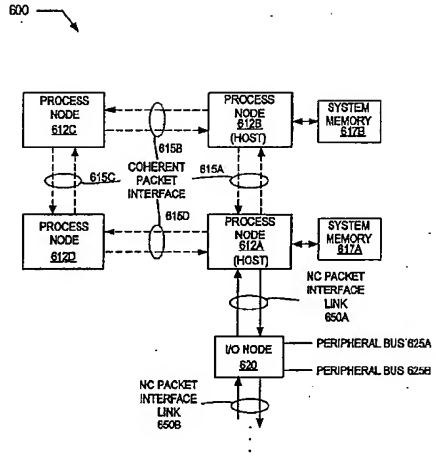

FIG. 6